

Superpolynomial speedups based on almost any quantum circuit

Sean Hallgren¹ and Aram W. Harrow²

¹ Department of Computer Science and Engineering, The Pennsylvania State University
University Park, PA

² Department of Mathematics, University of Bristol, Bristol, U.K.
a.harrow@bris.ac.uk

Abstract. The first separation between quantum polynomial time and classical bounded-error polynomial time was due to Bernstein and Vazirani in 1993. They first showed a $O(1)$ vs. $\Omega(n)$ quantum-classical oracle separation based on the quantum Hadamard transform, and then showed how to amplify this into a $n^{O(1)}$ time quantum algorithm and a $n^{\Omega(\log n)}$ classical query lower bound.

We generalize both aspects of this speedup. We show that a wide class of unitary circuits (which we call *dispersing* circuits) can be used in place of Hadamards to obtain a $O(1)$ vs. $\Omega(n)$ separation. The class of dispersing circuits includes all quantum Fourier transforms (including over nonabelian groups) as well as nearly all sufficiently long random circuits. Second, we give a general method for amplifying quantum-classical separations that allows us to achieve a $n^{O(1)}$ vs. $n^{\Omega(\log n)}$ separation from any dispersing circuit.

1 Background

Understanding the power of quantum computation relative to classical computation is a fundamental question. When we look at which problems can be solved in quantum but not classical polynomial time, we get a wide range: quantum simulation, factoring, approximating the Jones polynomial, Pell's equation, estimating Gauss sums, period-finding, group order-finding and even detecting some mildly non-abelian symmetries [Sho97, Hal07, Wat01, FIM⁺03, vDHI03]. However, when we look at what algorithmic tools exist on a quantum computer, the situation is not nearly as diverse. Apart from the BQP-complete problems [AJL06], the main tool for solving most of these problems is a quantum Fourier transform (QFT) over some group. Moreover, the successes have been for cases where the group is abelian or close to abelian in some way. For sufficiently nonabelian groups, there has been no indication that the transforms are useful even though they can be computed exponentially faster than classically. For example, while an efficient QFT for the symmetric group has been intensively studied for over a decade because of its connection to graph isomorphism, it is still unknown whether it can be used to achieve any kind of speedup over classical computation [Bea97].

The first separation between quantum computation and randomized computation was the Recursive Fourier Sampling problem (RFS) [BV97]. This algorithm had two components, namely using a Fourier transform, and using recursion. Shortly after this, Simon's algorithm and then Shor's algorithm for factoring were discovered, and the techniques from these algorithms have been the focus of most quantum algorithmic research since [Sim97, Sho97]. These developed into the hidden subgroup framework. The hidden subgroup problem is an oracle problem, but solving certain cases of it would result in solutions for factoring, graph isomorphism, and certain shortest lattice vector problems. Indeed, it was hoped that an algorithm for graph isomorphism could be found, but recent evidence suggests that this approach may not lead to one [HMR⁺06]. As a way to understand new techniques, this oracle problem has been very important, and it is also one of the very few where super-polynomial speedups have been found [IMS01, BCvD05].

In comparison to factoring, the RFS problem has received much less attention. The problem is defined as a property of a tree with labeled nodes and it was proven to be solvable with a quantum algorithm super-polynomially faster than the best randomized algorithm. This tree was defined in terms of the Fourier coefficients over \mathbb{Z}_2^n . The definition was rather technical, and it seemed that the simplicity of the Fourier coefficients for this group was necessary for the construction to work. Even the variants introduced by

Aaronson [Aar03] were still based on the same QFT over \mathbb{Z}_2^n , which seemed to indicate that this particular abelian QFT was a key part of the quantum advantage for RFS.

The main result of this paper is to show that the RFS structure can be generalized far more broadly. In particular, we show that an RFS-style super-polynomial speedup is achievable using almost any quantum circuit, and more specifically, it is also true for any Fourier transform (even nonabelian), not just over \mathbb{Z}_2^n . This illustrates a more general power that quantum computation has over classical computation when using recursion. The condition for a quantum circuit to be useful for an RFS-style speedup is that the circuit be *dispersing*, a concept we introduce to mean that it takes many different inputs to fairly even superpositions over most of the computational basis.

Our algorithm should be contrasted with the original RFS algorithm. One of the main differences between classical and quantum computing is so-called garbage that results from computing. It is important in certain cases, and crucial in recursion-based quantum algorithms because of quantum superpositions, that intermediate computations are uncomputed and that errors do not compound. The original RFS paper [BV97] avoided the error issue by using an oracle problem where every quantum state created from it had the exact property necessary with no errors. Their algorithm could have tolerated polynomially small errors, but in this paper we relax this significantly. We show that even if we can only create states with constant accuracy at each level of recursion, we can still carry through a recursive algorithm which introduces new constant-sized errors a polynomial number of times.

The main technical part of our paper shows that most quantum circuits can be used to construct separations relative to appropriate oracles. To understand the difficulty here, consider two problems that occur when one tries to define an oracle whose output is related to the amplitudes that result from running a circuit. First, it is not clear how to implement such an oracle since different amplitudes have different magnitudes, and only phases can be changed easily. Second, we need an oracle where we can prove that a classical algorithm requires many queries to solve the problem. If the oracle outputs many bits, this can be difficult or impossible to achieve. For example, the matrix entries of nonabelian groups can quickly reveal which representation is being used. To overcome these two problems we show that there are binary-valued functions that can approximate the complex-valued output of quantum circuits in a certain way.

One by-product of our algorithm is related to the Fourier transform of the symmetric group. Despite some initial promise for solving graph isomorphism, the symmetric group QFT has still not found any application in quantum algorithms. One instance of our result is the first example of a problem (albeit a rather artificial one) where the QFT over the symmetric group is used to achieve a super-polynomial speedup.

2 Statement of results

Our main contributions are to generalize the RFS algorithm of [BV97] in two stages. First, [BV97] described the problem of Fourier sampling over \mathbb{Z}_2^n , which has an $O(1)$ vs. $\Omega(n)$ separation between quantum and randomized complexities. We show that here the QFT over \mathbb{Z}_2^n can be replaced with a QFT over any group, or for that matter with almost any quantum circuit. Next, [BV97] turned Fourier sampling into recursive Fourier sampling with a recursive technique. We will generalize this construction to cope with error and to amplify a larger class of quantum speedups. As a result, we can turn any of the linear speedups we have found into superpolynomial speedups.

Let us now explain each of these steps in more detail. We replace the $O(1)$ vs $\Omega(n)$ separation based on Fourier sampling with a similar separation based on a more general problem called *oracle identification*. In the oracle identification problem, we are given access to an oracle $\mathcal{O}_a : X \rightarrow \{0, 1\}$ where $a \in A$, for some sets A and X with $\log |A|, \log |X| = \Theta(n)$. Our goal is to determine the identity of a . Further, assume that

we have access to a testing oracle $T_a : A \rightarrow \{0, 1\}$ defined by $T_a(a') = \delta_{a,a'}$, that will let us confirm that we have the right answer.³

A quantum algorithm for identifying a can be described as follows: first prepare a state $|\varphi_a\rangle$ using q queries to \mathcal{O}_a , then perform a POVM $\{\Pi_{a'}\}_{a' \in A}$ (with $\sum_{a'} \Pi_{a'} \leq I$ to allow for the possibility of a “failure” outcome), using no further queries to \mathcal{O}_a . The success probability is $\langle \varphi_a | \Pi_a | \varphi_a \rangle$. For our purposes, it will suffice to place a $\Omega(1)$ lower bound on this probability: say that for each a , $\langle \varphi_a | \Pi_a | \varphi_a \rangle \geq \delta$ for some constant $\delta > 0$. On the other hand, any classical algorithm trivially requires $\geq \log(|A|\delta) = \Omega(n)$ oracle calls to identify a with success probability $\geq \delta$. This is because each query returns only one bit of information. In Theorem 9 we will describe how a large class of quantum circuits can achieve this $O(1)$ vs. $\Omega(n)$ separation, and in Theorems 11 and 12 we will show specifically that QFTs and most random circuits fall within this class.

Now we describe the amplification step. This is a variant of the [BV97] procedure in which making an oracle call in the original problem requires solving a sub-problem from the same family as the original problem. Iterating this ℓ times turns query complexity q into $q^{\Theta(\ell)}$, so choosing $\ell = \Theta(\log n)$ will yield the desired polynomial vs. super-polynomial separation. We will generalize this construction by defining an amplified version of oracle identification called *recursive oracle identification*. This is described in the next section, where we will see how it gives rise to superpolynomial speedups from a broad class of circuits.

We conclude that quantum speedups—even superpolynomial speedups—are much more common than the conventional wisdom would suggest. Moreover, as useful as the QFT has been to quantum algorithms, it is far from the only source of quantum algorithmic advantage.

3 Recursive amplification

In this section we show that once we are given a constant versus linear separation (for quantum versus classical oracle identification), we are able to amplify this to a super-polynomial speedup. We require a much looser definition than in [BV97] because the constant case can have a large error.

Definition 1. For sets A, X , let $f : A \times X \rightarrow \{0, 1\}$ be a function. To set the scale of the problem, let $|X| = 2^n$ and $|A| = 2^{\Omega(n)}$. Define the set of oracles $\{\mathcal{O}_a : a \in A\}$ by $\mathcal{O}_a(x) = f(a, x)$, and the states $|\varphi_a\rangle = \frac{1}{\sqrt{|X|}} \sum_{x \in X} (-1)^{f(a,x)} |x\rangle$. The single-level oracle identification problem is defined to be the task of determining a given access to \mathcal{O}_a . Let U be a family of quantum circuits, implicitly depending on n . We say that U solves the single-level oracle identification problem if

$$|\langle a | U | \varphi_a \rangle|^2 \geq \Omega(1)$$

for all sufficiently large n and all $a \in A$. In this case, we define the POVM $\{\Pi_a\}_{a \in A}$ by $\Pi_a = U^\dagger |a\rangle \langle a| U$.

When this occurs, it means that a can be identified from \mathcal{O}_a with $\Omega(1)$ success probability and using a single query. In the next section, we will show how a broad class of unitaries U (the so-called *dispersing* unitaries) allow us to construct f for which U solves the single-level oracle identification problem. There are natural generalizations to oracle identification problems requiring many queries, but we will not explore them here.

Theorem 2. Suppose we are given a single-level oracle problem with function f and unitary U running in time $\text{poly}(n)$. Then we can construct a modified oracle problem from f which can be solved by a quantum

³ This will later allow us to turn two-sided into one-sided error; unfortunately it also means that a non-deterministic Turing machine can find a with a single query to T_a . Thus, while the oracle defined in BV is a candidate for placing BQP outside PH, ours will not be able to place BQP outside of NP. This limitation appears not to be fundamental, but we will leave the problem of circumventing it to future work.

computer in polynomial time (and queries), but requires $n^{\Omega(\log n)}$ queries for any classical algorithm that succeeds with probability $\frac{1}{2} + n^{-o(\log n)}$.

We start by defining the modified version of the problem (Definition 3 below), and describing a quantum algorithm to solve it. Then in Theorem 4 we will show that the quantum algorithm solves the problem correctly in polynomial time, and in Theorem 6, we will show that randomized classical algorithms require superpolynomial time to have a nonnegligible probability of success.

The recursive version of the problem simply requires that another instance of the problem be solved in order to access a value at a child. Figure 1 illustrates the structure of the problem.

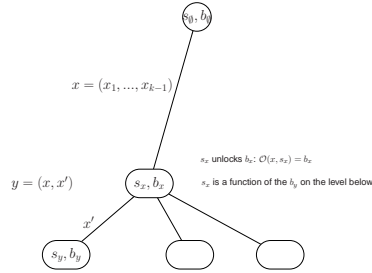


Fig. 1. A depth k node at location $x = (x_1, \dots, x_k)$ is labeled by its secret s_x and a bit b_x . The secret s_x can be computed from the bits b_y of its children, and once it is known, the bit b_x is computed from the oracle $\mathcal{O}(x, s_x) = b_x$. If x is a leaf then it has no secret and we simply have $b_x = \mathcal{O}(x)$. The goal is to compute the secret bit b_0 at the root.

Using the notation from Figure 1, the relation between a secret s_x , and the bits b_y of its children is given by $b_y = f(s_x, x')$, where f is the function from the single-level oracle identification problem. Thus by computing enough of the bits b_{y_1}, b_{y_2}, \dots corresponding to children y_1, y_2, \dots , we can solve the single-level oracle identification problem to find s_x . Of course computing the b_y will require finding the secret strings s_y , which requires finding the bits of *their* children and so on, until we reach the bottom layer where queries return answer bits without the need to first produce secret strings.

Definition 3. A level- ℓ recursive oracle identification problem is specified by X, A and f from a single-level oracle identification problem (Definition 1), any function $s : \emptyset \cup X \cup X \times X \cup \dots \cup X^{\ell-1} \rightarrow A$, and any final answer $b_0 \in \{0, 1\}$. Given these ingredients, an oracle \mathcal{O} is defined which takes inputs in

$$\bigcup_{k=0}^{\ell-1} [X^k \times A] \cup X^\ell$$

and to return outputs in $\{0, 1, \text{FAIL}\}$. On inputs $x_1, \dots, x_k \in X, a \in A$ with $1 \leq k < \ell$, \mathcal{O} returns

$$\mathcal{O}(x_1, \dots, x_k, a) = f(s(x_1, \dots, x_{k-1}), x_k) \quad \text{when } a = s(x_1, \dots, x_k) \quad (1)$$

$$\mathcal{O}(x_1, \dots, x_k, a) = \text{FAIL} \quad \text{when } a \neq s(x_1, \dots, x_k). \quad (2)$$

If $k = 0$, then $\mathcal{O}(s(\emptyset)) = b_\emptyset$ and $\mathcal{O}(a) = \text{FAIL}$ if $a \neq s(\emptyset)$. When $k = \ell$,

$$\mathcal{O}(x_1, \dots, x_\ell) = f(s(x_1, \dots, x_{\ell-1}), x_\ell).$$

The recursive oracle identification problem is to determine b_\emptyset given access to \mathcal{O} .

Note that the function s gives the values s_x in Figure 1. These values are actually defined in the oracle and can be chosen arbitrarily at each node. Note also that the oracle defined here effectively includes a testing oracle, which can determine whether $a = s(x_1, \dots, x_k)$ for any $a \in A, x_1, \dots, x_k \in X$ with one query. (When $x = (x_1, \dots, x_k)$, we use $s(x_1, \dots, x_k)$ and s_x interchangeably.) A significant difference between our construction and that of [BV97] is that the values of s at different nodes can be set completely independently in our construction, whereas [BV97] had a complicated consistency requirement.

The algorithm. Now we turn to a quantum algorithm for the recursive oracle identification problem. If a quantum computer can identify a with one-sided⁴ error $1 - \delta$ using time T and q queries in the non-recursive problem, then we will show that the recursive version can be solved in time $O((q \frac{\log 1/\delta}{\delta})^\ell T)$. For concreteness, suppose that $|\varphi_a\rangle = \frac{1}{\sqrt{|X|}} \sum_{x \in X} (-1)^{f(a,x)} |x\rangle$, so that $q = 1$; the case when $q > 1$ is an easy, but tedious, generalization. Suppose that our identifying quantum circuit is U , so a can be identified by applying the POVM $\{\Pi_{a'}\}_{a' \in A}$ with $\Pi_{a'} = U^\dagger |a'\rangle\langle a'| U$ to the state $|\varphi_a\rangle$.

The intuitive idea behind our algorithm is as follows: At each level, we find $s(x_1, \dots, x_k)$ by recursively computing $s(x_1, \dots, x_{k+1})$ for each x_{k+1} (in superposition) and using this information to create many copies of $|\varphi_{s(x_1, \dots, x_k)}\rangle$, from which we can extract our answer. However, we need to account for the errors carefully so that they do not blow up as we iterate the recursion. In what follows, we will adopt the convention that Latin letters in kets (e.g. $|a\rangle, |x\rangle, \dots$) denote computational basis states, while Greek letters (e.g. $|\zeta\rangle, |\varphi\rangle, \dots$) are general states that are possibly superpositions over many computational basis states. Also, we let the subscript (k) indicate a dependence on (x_1, \dots, x_k) . The recursive oracle identification algorithm is as follows:

Algorithm: FIND

Input: $|x_1, \dots, x_k\rangle|0\rangle$ for $k < \ell$

Output: $a_{(k)} = s(x_1, \dots, x_k)$ up to error $\varepsilon = (\delta/8)^2$, where δ is the constant from the oracle. This means $|x_1, \dots, x_k\rangle \left[\sqrt{1 - \varepsilon_{(k)}} |0\rangle |a_{(k)}\rangle |\zeta_{(k)}\rangle + \sqrt{\varepsilon_{(k)}} |1\rangle |\zeta'_{(k)}\rangle \right]$, where $\varepsilon_{(k)} \leq \varepsilon$ and $|\zeta_{(k)}\rangle$ and $|\zeta'_{(k)}\rangle$ are arbitrary. (We can assume this form without loss of generality by absorbing phases into $|\zeta_{(k)}\rangle$ and $|\zeta'_{(k)}\rangle$.)

1. Create the superposition $\frac{1}{\sqrt{|X|}} \sum_{x_{k+1} \in X} |x_{k+1}\rangle$.
2. If $k + 1 < \ell$ then let $a_{(k+1)} = \text{FIND}(x_1, \dots, x_{k+1})$ (with error $\leq \varepsilon$), otherwise $a_{(k+1)} = \emptyset$.
3. Call the oracle $\mathcal{O}(x_1, \dots, x_{k+1}, a_{(k+1)})$ to apply the phase $(-1)^{f(s(x_1, \dots, x_k), x_{k+1})}$ using the key $a_{(k+1)}$.
4. If $k + 1 < \ell$ then call FIND^\dagger to (approximately) uncompute $a_{(k+1)}$.
5. We are now left with $|\tilde{\varphi}_{(k)}\rangle$, which is close to $|\varphi_{s(x_1, \dots, x_k)}\rangle$. Repeat steps 1–4 $m = \frac{4}{\delta} \ln \frac{8}{\delta}$ times to obtain $|\tilde{\varphi}_{(k)}\rangle^{\otimes m}$.
6. Coherently measure $\{\Pi_a\}$ on each copy and test the results (i.e. apply U , test the result, and apply U^\dagger).
7. If any tests pass, copy the correct $a_{(k)}$ to an output register, along with $|0\rangle$ to indicate success.

⁴ One-sided error is a reasonable demand given our access to a testing oracle. Most of these results go through with two-sided error as well, but for notational simplicity, we will not explore them here.

Otherwise put a $|1\rangle$ in the output to indicate failure.

8. Let everything else comprise the junk register $|\zeta_{(k)}\rangle$.

Theorem 4. *Calling FIND on $|0\rangle$ solves the recursive oracle problem in quantum polynomial time.*

Proof. The proof is by backward induction on k ; we assume that the algorithm returns with error $\leq \varepsilon$ for $k+1$ and prove it for k . The initial step when $k = \ell$ is trivial since there is no need to compute $a_{\ell+1}$, and thus no source of error. If $k < \ell$, then assume that correctness of the algorithm has already been proved for $k+1$. Therefore Step 2 leaves the state

$$\frac{1}{\sqrt{|X|}} \sum_{x_{k+1} \in X} |x_{k+1}\rangle \left[\sqrt{1 - \varepsilon_{(k+1)}} |0\rangle |a_{(k+1)}\rangle |\zeta_{(k+1)}\rangle + \sqrt{\varepsilon_{(k+1)}} |1\rangle |\zeta'_{(k+1)}\rangle \right].$$

In Step 3, we assume for simplicity that the oracle was called conditional on the success of Step 2. This yields

$$|\psi'_{(k)}\rangle := \frac{1}{\sqrt{|X|}} \sum_{x_{k+1} \in X} |x_{k+1}\rangle \left[(-1)^{f(a_{(k)}, x_{k+1})} \sqrt{1 - \varepsilon_{(k+1)}} |0\rangle |a_{(k+1)}\rangle |\zeta_{(k+1)}\rangle + \sqrt{\varepsilon_{(k+1)}} |1\rangle |\zeta'_{(k+1)}\rangle \right].$$

Now define the state $|\psi_{(k)}\rangle$ by

$$|\psi_{(k)}\rangle := \frac{1}{\sqrt{|X|}} \sum_{x_{k+1} \in X} (-1)^{f(a_{(k)}, x_{k+1})} |x_{k+1}\rangle \left[\sqrt{1 - \varepsilon_{(k+1)}} |0\rangle |a_{(k+1)}\rangle |\zeta_{(k+1)}\rangle + \sqrt{\varepsilon_{(k+1)}} |1\rangle |\zeta'_{(k+1)}\rangle \right].$$

Note that

$$\langle \psi'_{(k)} | \psi_{(k)} \rangle = \frac{1}{|X|} \sum_{x_{k+1} \in X} \left(1 - \varepsilon_{(k+1)} + (-1)^{f(a_{(k)}, x_{k+1})} \varepsilon_{(k+1)} \right).$$

This quantity is real and always $\geq 1 - 2\varepsilon_{(k+1)} \geq \sqrt{1 - 4\varepsilon}$ by the induction hypothesis. Let

$$|\phi_{(k)}\rangle := \frac{1}{|X|} \sum_{x_{k+1} \in X} (-1)^{f(a_{(k)}, x_{k+1})} |x_{k+1}\rangle |0\rangle.$$

Note that $\text{FIND}^\dagger |x_1, \dots, x_k, \psi_{(k)}\rangle = |x_1, \dots, x_k, \phi_{(k)}\rangle$. Thus there exists $\varepsilon_{(k)}$ such that applying FIND^\dagger to $|x_1, \dots, x_k\rangle |\psi'_{(k)}\rangle$ yields

$$|x_1, \dots, x_k\rangle \otimes \left[\sqrt{1 - 4\varepsilon_{(k)}} |\phi_{(k)}\rangle + \sqrt{4\varepsilon_{(k)}} |\phi'_{(k)}\rangle \right],$$

where $\langle \phi_{(k)} | \phi'_{(k)} \rangle = 0$ and $\varepsilon_{(k)} \leq \varepsilon$.

We now want to analyze the effects of measuring $\{\Pi_a\}$ when we are given the state

$$|\varphi_{(k)}\rangle := \sqrt{1 - 4\varepsilon_{(k)}} |\phi_{(k)}\rangle + \sqrt{4\varepsilon_{(k)}} |\phi'_{(k)}\rangle$$

instead of $|\phi_{(k)}\rangle$. If we define $\|M\|_1 = \text{tr} \sqrt{M^\dagger M}$ for a matrix M , then $\| |\varphi_{(k)}\rangle \langle \varphi_{(k)}| - |\phi_{(k)}\rangle \langle \phi_{(k)}| \|_1 = 4\sqrt{\varepsilon_{(k)}}$ [FvdG99]. Thus

$$\langle \varphi_{(k)} | \Pi_{a_{(k)}} | \varphi_{(k)} \rangle \geq \langle \phi_{(k)} | \Pi_{a_{(k)}} | \phi_{(k)} \rangle - 4\sqrt{\varepsilon_{(k)}} \geq \delta - 4\sqrt{\varepsilon_{(k)}} \geq \delta/2.$$

In the last step we have chosen $\varepsilon = (\delta/8)^2$.

Finally, we need to guarantee that with probability $\geq 1 - \varepsilon$ at least one of the tests in Step 6 passes. After applying U and the test oracle to $|\varphi_{(k)}\rangle$, we have $\geq \sqrt{\delta/2}$ overlap with a successful test and $\leq \sqrt{1 - \delta/2}$ overlap with an unsuccessful test. When we repeat this m times, the amplitude in the subspace corresponding to all tests failing is $\leq (1 - \delta/2)^{m/2} \leq e^{-m\delta/4}$. If we choose $m = (2/\delta) \ln(1/\varepsilon) = (4/\delta) \ln(8/\delta)$ then the failure amplitude will be $\leq \sqrt{\varepsilon}$, as desired.

To analyze the time complexity, first note that the run-time is $O(T)$ times the number of queries made by the algorithm, and we have assumed that T is polynomial in n . Suppose the algorithm at level k requires $Q(k)$ queries. Then steps 2 and 4 require $mQ(k+1)$ queries each, steps 3 and 6 require m queries each and together $Q(k) = 2mQ(k+1) + 2m$. The base case is $k = \ell$, for which $Q(\ell) = 0$, since there are no secret strings to calculate for the leaves. The total number of queries required for the algorithm is then $Q(0) \approx (2m)^{2^\ell}$. If we choose $\ell = \log n$ the quantum query complexity will thus be $n^{2 \log 2m} = n^{O(1)}$ and the quantum complexity will be polynomial in n compared with the $n^{\Omega(\log n)}$ lower bound.

This concludes the demonstration of the polynomial-time quantum algorithm. Now we turn to the classical $n^{\Omega(\log n)}$ lower bound. Our key technical result is the following lemma:

Lemma 5. *Define the recursive oracle identification problem as above, with a function $f : A \times X \rightarrow \{0, 1\}$ and a secret $s : \emptyset \cup X \cup X \times X \cup \dots \cup X^{\ell-1} \mapsto A$ encoded in an oracle \mathcal{O} . Fix a deterministic classical algorithm that makes $\leq Q$ queries to \mathcal{O} . Then if s and ANS are chosen uniformly at random, the probability that ANS is output by the algorithm is*

$$\leq \frac{1}{2} + \max \left(\frac{Q}{|A|^{1/3} - Q}, Q \left(\frac{\log |A|}{3} \right)^{-\ell} \right).$$

Using Yao's minimax principle and plugging in $|A| = 2^{\alpha n}$, $\ell = \log n$ and $Q = n^{o(\log n)}$ readily yields

Theorem 6. *If $\log |A| = n^{\Omega(1)}$ and $\ell = \Omega(\log n)$, then any randomized classical algorithm using $Q = n^{o(\log n)}$ queries will have $\frac{1}{2} + n^{-\Omega(\log n)}$ probability of successfully outputting ANS.*

Proof (of Lemma 5). Let $T = \emptyset \cup X \cup \dots \cup X^\ell$ denote the tree on which the oracle is defined. We say that a node $x \in T$ has been *hit* by the algorithm if position x has been queried by the oracle together with the correct secret, i.e. $\mathcal{O}(s(x), x)$ has been queried. The only way to find to obtain information about ANS is for the algorithm to query \emptyset with the appropriate secret; in other words, to hit \emptyset .

For $x, y \in T$ we say that x is an *ancestor* of y , and that y is a *descendant* of x , if $y = x \times z$ for some $z \in T$. If $z \in X$ then we say that y is a *child* of x and that x is a *parent* of y . Now define $S \subset T$ to be the set of all $x \in T$ such that x has been hit but none of x 's ancestors have been. Also define a function $d(x)$ to be the depth of a node x ; i.e. for all $x \in X^k$, $d(x) = k$. We combine these definitions to declare an invariant

$$Z = \sum_{x \in S} \left(\frac{\log |A|}{3} \right)^{-d(x)}$$

The key properties of Z we need are that:

1. Initially $Z = 0$.
2. If the algorithm is successful then it terminates with $Z = 1$.
3. Only oracle queries change the value of Z .
4. Querying a leaf can add at most $(\log |A|/3)^{-\ell}$ to Z .
5. Querying an internal node (i.e. not a leaf) can add at most $2/(|A|^{1/3} - Q)$ to $\mathbf{E} Z$, where \mathbf{E} indicates the expectation over random choices of s .

Combining these facts yields the desired bound.

Properties 1–4 follow directly from the definition (with the inequality in property 4 because it is possible to query a node that has already been hit). To establish property 5, suppose that the algorithm queries node $x \in T$ and that it has previously hit k of x 's children. This gives us some partial information about $s(x)$. We can model this information as a partition of A into 2^k disjoint sets A_1, \dots, A_{2^k} (of which some could be empty). From the k bits returned by the oracle on the k children of x we have successfully queried, we know not only that $s(x) \in A$, but that $s(x) \in A_i$ for some $i \in \{1, \dots, 2^k\}$.

We will now divide the analysis into two cases. Either $k \leq \frac{1}{3} \log |A|$ or $k > \frac{1}{3} \log |A|$. We will argue that in the former case, $|A_i|$ is likely to be large, and so we are unlikely to successfully guess $s(x)$, while in the latter case even a successful guess will not increase Z . The latter case ($k > \frac{1}{3} \log |A|$) is easier, so we consider it first. In this case, Z only changes if x is hit in this step and neither x nor any of its ancestors have been previously hit. Then even though hitting x will contribute $(\log |A|/3)^{-d(x)}$ to Z , it will also remove the k children from S (as well as any other descendants of x), which will decrease Z by at least $k(\log |A|/3)^{-d(x)-1} > (\log |A|/3)^{-d(x)}$, resulting in a net decrease of Z .

Now suppose that $k \leq \frac{1}{3} \log |A|$. Recall that our information about $s(x)$ can be expressed by the fact that $s(x) \in A_i$ for some $i \in \{1, \dots, 2^k\}$. Since the values of s were chosen uniformly at random, we have $\Pr(A_i) = |A_i|/|A|$. Say that a set A_i is *bad* if $|A_i| \leq |A|^{2/3}/2^k$. Then for a particular bad set A_i , $\Pr(A_i) \leq |A|^{-1/3}2^{-k}$. From the union bound, we see that the probability that *any* bad set is chosen is $\leq |A|^{-1/3}$.

Assume then that we have chosen a good set A_i , meaning that conditioned on the values of the children there are $|A_i| \geq |A|^{2/3}/2^k \geq |A|^{1/3}$ possible values of $s(x)$. However, previous failed queries at x may also have ruled out specific possible values of x . There have been at most Q queries at x , so there are $\geq |A|^{1/3} - Q$ possible values of $s(x)$ remaining. (Queries to any other nodes in the graph yield no information on $s(x)$.) Thus the probability of hitting x is $\leq 1/(|A|^{1/3} - Q)$ if we have chosen a good set. We also have a $\leq |A|^{-1/3}$ probability of choosing a bad set, so the total probability of hitting x (in the $k \leq \frac{1}{3} \log |A|$ case) is $\leq |A|^{-1/3} + 1/(|A|^{1/3} - Q) \leq 2/(|A|^{1/3} - Q)$. Finally, hitting x will increase Z by at most one, so the largest possible increase of $\mathbb{E} Z$ when querying a non-leaf node is $\leq 2/(|A|^{1/3} - Q)$. This completes the proof of property 5 and thus the Lemma.

4 Dispersing Circuits

In this section we define *dispersing* circuits and show how to construct an oracle problem with a constant versus linear separation from any such circuit. In the next sections we will show how to find dispersing circuits. Our strategy for finding speedups will be to start with a unitary circuit U which acts on n qubits and has size polynomial in n . We will then try to find an oracle for which U efficiently solves the corresponding oracle identification problem. Next we need to define a state $|\varphi_a\rangle$ that can be prepared with $O(1)$ oracle calls and has $\Omega(1)$ overlap with $U^\dagger|a\rangle$. This is accomplished by letting $|\varphi_a\rangle$ be a state of the form $2^{-n/2} \sum_x \pm |x\rangle$. We can prepare $|\varphi_a\rangle$ with only two oracle calls (or one, depending on the model), but to guarantee that $|\langle a|U|\varphi_a\rangle|$ can be made large, we will need an additional condition on U . For any $a \in A$, $U^\dagger|a\rangle$ should have amplitude that is mostly spread out over the entire computational basis. When this is the case, we say that U is *dispersing*. The precise definition is as follows:

Definition 7. Let U be a quantum circuit on n qubits. For $0 < \alpha, \beta \leq 1$, we say that U is (α, β) -dispersing if there exists a set $A \subseteq \{0, 1\}^n$ with $|A| \geq 2^{\alpha n}$ and

$$\sum_{x \in \{0, 1\}^n} |\langle a|U|x\rangle| \geq \beta 2^{\frac{n}{2}}. \quad (3)$$

for all $a \in A$.

Note that the LHS of (3) can also be interpreted as the L_1 norm of $U^\dagger|a\rangle$.

The speedup in [BV97] uses $U = H^{\otimes n}$, which is (1,1)-dispersing since $\sum_x |\langle a|H^{\otimes n}|x\rangle| = 2^{n/2}$ for all a . Similarly the QFT over the cyclic group is (1,1)-dispersing.⁵ Nonabelian QFTs do not necessarily have the same strong dispersing properties, but they satisfy a weaker definition that is still sufficient for a quantum speedup. Suppose that the measurement operator is instead defined as $\Pi_a = U(|a\rangle\langle a| \otimes I)U^\dagger$, where a is a string on m bits and I denotes the identity operator on $n - m$ bits. Then U still permits oracle identification, but our requirements that U be dispersing are now relaxed. Here, we give a definition that is loose enough for our purposes, although further weakening would still be possible.

Definition 8. Let U be a quantum circuit on n qubits. For $0 < \alpha, \beta \leq 1$ and $0 < m \leq n$, we say that U is (α, β) -pseudo-dispersing if there exists a set $A \subseteq \{0, 1\}^m$ with $|A| \geq 2^{\alpha n}$ such that for all $a \in A$ there exists a unit vector $|\psi\rangle \in \mathbb{C}^{2^{n-m}}$ such that

$$\sum_{x \in \{0, 1\}^n} |\langle a|\langle \psi|U|x\rangle| \geq \beta 2^{\frac{n}{2}}. \quad (4)$$

This is a weaker property than being dispersing, meaning that any (α, β) -dispersing circuit is also (α, β) -pseudo-dispersing.

We can now state our basic constant vs. linear query separation.

Theorem 9. If U is (α, β) -pseudo-dispersing, then there exists an oracle problem which can be solved with one query, one use of U and success probability $(2\beta/\pi)^2$. However, any classical randomized algorithm that succeeds with probability $\geq \delta$ must use $\geq \alpha n + \log \delta$ queries.

Before we prove this Theorem, we state a Lemma about how well states of the form $2^{-n/2} \sum_x e^{i\phi_x} |x\rangle$ can be approximated by states of the form $2^{-n/2} \sum_x \pm |x\rangle$.

Lemma 10. For any vector $(x_1, \dots, x_d) \in \mathbb{C}^d$ there exists $(\theta_1, \dots, \theta_d) \in \{\pm 1\}^d$ such that

$$\left| \sum_{k=1}^d x_k \theta_k \right| \geq \frac{2}{\pi} \sum_{k=1}^d |x_k|.$$

The proof is in the full version of the paper[?].

Proof of Theorem 9: Since U is (α, β) -pseudo-dispersing, there exists a set $A \subset \{0, 1\}^m$ with $|A| \geq 2^{\alpha n}$ and satisfying (4) for each $a \in A$. The problem will be to determine a by querying an oracle $\mathcal{O}_a(x)$. No matter how we define the oracle, as long as it returns only one bit per call any classical randomized algorithm making q queries can have success probability no greater than $2^{q-\alpha n}$ (or else guessing could succeed with probability $> 2^{-\alpha n}$ without making any queries). This implies the classical lower bound.

Given $a \in A$, to define the oracle \mathcal{O}_a , first use the definition to choose a state $|\psi\rangle$ satisfying (4). Then by Lemma 10 (below), choose a vector θ that (when normalized to $|\theta\rangle$) will approximate the state $U^\dagger|a\rangle|\psi\rangle$. Define $\mathcal{O}_a(x)$ so that $(-1)^{\mathcal{O}_a(x)} = \theta_x = 2^{n/2} \langle x|\theta\rangle$. By construction,

$$2^{-n/2} |\langle a|\langle \psi|U|\theta\rangle| \geq \frac{2}{\pi} \beta \quad (5)$$

which implies that creating $|\theta\rangle$, applying U , and measuring the first register has probability $\geq (2\beta/\pi)^2$ of yielding the correct answer a . \square

⁵ Another possible way to generalize [BV97] is to consider other unitaries of the form $U = A^{\otimes n}$, for $A \in \mathcal{U}_2$. However, it is not hard to show that the only way for such a U to be $(\Omega(1), \Omega(1))$ -dispersing is for A to be of the form $e^{i\phi_1 \sigma_z} H e^{i\phi_2 \sigma_z}$.

5 Any quantum Fourier transform is pseudo-dispersing

In this section we start with some special cases of dispersing circuits by showing that any Fourier transform is dispersing. In the next section we show that most circuits are dispersing.

The original RFS paper [BV97] used the fact that $H^{\otimes n}$ is (1,1)-dispersing to obtain their starting $O(1)$ vs $\Omega(n)$ separation. The QFT on the cyclic group (or any abelian group, in fact) is also (1,1)-dispersing. In fact, if we will accept a pseudo-dispersing circuit, then any QFT will work:

Theorem 11. *Let G be a group with irreps \hat{G} and d_λ denoting the dimension of irrep λ . Then the Fourier transform over G is $(\alpha, 1/\sqrt{2})$ -pseudo-dispersing, where $\alpha = (\log \sum_\lambda d_\lambda) / \log |G| \geq 1/2$.*

Via Theorem 9 and Theorem 2, this implies that any QFT can be used to obtain a superpolynomial quantum speedup. For most nonabelian QFTs, this is the first example of a problem which they can solve more quickly than a classical computer.

Proof (Proof of Theorem 11). Let $A = \{(\lambda, i) : \lambda \in \hat{G}, i \in \{1, \dots, d_\lambda\}\}$.

Let V_λ denote the representation space corresponding to an irrep $\lambda \in \hat{G}$. The Fourier transform on G maps vectors in $\mathbb{C}[G]$ to superpositions of vectors of the form $|\lambda\rangle|v_1\rangle|v_2\rangle$ for $|v_1\rangle, |v_2\rangle \in V_\lambda$.

Fix a particular choice of λ and $|i\rangle \in V_\lambda$. If U denotes the QFT on G then let

$$\rho = U^\dagger \left(|\lambda\rangle\langle\lambda| \otimes |i\rangle\langle i| \otimes \frac{I_{V_\lambda}}{d_\lambda} \right) U.$$

Define $V := \text{supp } \rho$, and let $\mathbf{E}_{|\psi\rangle \in V}$ denote an expectation over $|\psi\rangle$ chosen uniformly at random from unit vectors in V .⁶ Finally, let Π be the projector onto V . Note that $\rho = \Pi/d_\lambda = \mathbf{E}_{|\psi\rangle \in V} |\psi\rangle\langle\psi|$.

Because of the invariance of ρ under right-multiplication by group elements (i.e. $\langle g_1|\rho|g_2\rangle = \langle g_1 h|\rho|g_2 h\rangle$ for all $g_1, g_2, h \in G$), we have for any g that

$$\langle g|\rho|g\rangle = \frac{1}{|G|} \sum_h \langle gh|\rho|gh\rangle = \frac{1}{|G|} \text{tr}(\rho) = \frac{1}{|G|}. \quad (6)$$

Since $\mathbf{E}_{|\psi\rangle \in V} |\psi\rangle\langle\psi| = \rho$, (6) implies that

$$\mathbf{E}_{|\psi\rangle \in V} |\langle g|\psi\rangle|^2 = \langle g|\rho|g\rangle = \frac{1}{|G|}.$$

Next, we would like to analyze $\mathbf{E}_{|\psi\rangle \in V} |\langle g|\psi\rangle|^4$.

$$\mathbf{E}_{|\psi\rangle} |\langle g|\psi\rangle|^4 = \mathbf{E}_{|\psi\rangle} \text{tr}(|g\rangle\langle g| \otimes |g\rangle\langle g| \cdot (|\psi\rangle\langle\psi| \otimes |\psi\rangle\langle\psi|)) \quad (7)$$

$$= \text{tr}(|g\rangle\langle g| \otimes |g\rangle\langle g|) \frac{I + \text{SWAP}}{d_\lambda(d_\lambda + 1)} (\Pi \otimes \Pi) \quad (8)$$

$$\leq \text{tr}(|g\rangle\langle g| \otimes |g\rangle\langle g|) \cdot (I + \text{SWAP})(\rho \otimes \rho) \quad (9)$$

$$= 2(\langle g|\rho|g\rangle)^2 = \frac{2}{|G|^2} \quad (10)$$

To prove the equality on the second line, we use a standard representation-theoretic trick (cf. section V.B of [PSW06]). First note that $|\psi\rangle^{\otimes 2}$ belongs to the symmetric subspace of $V \otimes V$, which is a $\frac{d_\lambda(d_\lambda+1)}{2}$ -dimensional irrep of \mathcal{U}_{d_λ} . Since $\mathbf{E}_{|\psi\rangle} |\psi\rangle\langle\psi|^{\otimes 2}$ is invariant under conjugation by $u \otimes u$ for any $u \in \mathcal{U}_{d_\lambda}$,

⁶ We can think of $|\psi\rangle$ either as the result of applying a Haar uniform unitary to a fixed unit vector, or by choosing $|\psi'\rangle$ from any rotationally invariant ensemble (e.g. choosing the real and imaginary part of each component to be an i.i.d. Gaussian with mean zero) and setting $|\psi\rangle = |\psi'\rangle / \sqrt{\langle\psi'|\psi'\rangle}$.

it follows that $\mathbf{E}_{|\psi\rangle} |\psi\rangle\langle\psi|^{\otimes 2}$ is proportional to a projector onto the symmetric subspace of $V^{\otimes 2}$. Finally, $\text{SWAP}\Pi^{\otimes 2}$ has eigenvalue 1 on the symmetric subspace of $V^{\otimes 2}$ and eigenvalue -1 on its orthogonal complement, the antisymmetric subspace of $V^{\otimes 2}$. Thus, $\frac{I+\text{SWAP}}{2}\Pi^{\otimes 2}$ projects onto the symmetric subspace and we conclude that

$$\mathbf{E}_{|\psi\rangle} |\psi\rangle\langle\psi|^{\otimes 2} = \frac{(I + \text{SWAP})(\Pi \otimes \Pi)}{d_\lambda(d_\lambda + 1)}.$$

Now we note the inequality

$$\mathbf{E} |Y| \geq (\mathbf{E} Y^2)^{\frac{3}{2}} / (\mathbf{E} Y^4)^{\frac{1}{2}}, \quad (11)$$

which holds for any random variable Y and can be proved using Hölder's inequality [Ber97]. Setting $Y = |\langle g|\psi\rangle|$, we can bound $\mathbf{E}_{|\psi\rangle} |\langle g|\psi\rangle| \geq 1/\sqrt{2|G|}$. Summing over G , we find

$$\mathbf{E}_{|\psi\rangle} \sum_{g \in G} |\langle g|\psi\rangle| \geq \frac{1}{\sqrt{2}} \sqrt{|G|}.$$

Finally, because this last inequality holds in expectation, it must also hold for at least some choice of $|\psi\rangle$. Thus there exists $|\psi\rangle \in V$ such that

$$\sum_{g \in G} |\langle g|\psi\rangle| \geq \frac{1}{\sqrt{2}} \sqrt{|G|}.$$

Then U satisfies the pseudo-dispersing condition in (4) for the state $|\psi\rangle$ with $\beta = 1/\sqrt{2}$.

This construction works for each $\lambda \in \hat{G}$ and for $|v_1\rangle$ running over any choice of basis of V_λ . Together, this comprises $\sum_{\lambda \in \hat{G}} d_\lambda$ vectors in the set A .

6 Most circuits are dispersing

Our final, and most general, method of constructing dispersing circuits is simply to choose a polynomial-size random circuit. We define a length- t random circuit to consist of performing the following steps t times.

1. Choose two distinct qubits i, j at random from $[n]$.
2. Choose a Haar-distributed random $U \in \mathcal{U}_4$.
3. Apply U to qubits i and j .

A similar model of random circuits was considered in [DOP07]. Our main result about these random circuits is the following Theorem.

Theorem 12. *For any $\alpha, \beta > 0$, there exists a constant C such that if U is a random circuit on n qubits of length $t = Cn^3$ then U is (α, β) -dispersing with probability*

$$\geq 1 - \frac{2\beta^2}{1 - 2^{-n(1-\alpha)}}.$$

Theorem 12 is proved in the extended version of this paper[?]. The idea of the proof is to reduce the evolution of the fourth moments of the random circuit (i.e. quantities of the form $\mathbf{E}_U \text{tr} U M_1 U^\dagger M_2 U M_3 U^\dagger M_4$) to a classical Markov chain, using the approach of [DOP07]. Then we show that this Markov chain has a gap of $\Omega(1/n^2)$, so that circuits of length $O(n^3)$ have fourth moments nearly identical to those of Haar-uniform unitaries from \mathcal{U}_{2^n} . Finally, we use (11), just as we did for quantum Fourier transforms, to show that a large fraction of inputs are likely to be mapped to states with large L_1 -norm. This will prove Theorem 12 and show that superpolynomial quantum speedups can be built by plugging almost any circuit into the recursive framework we describe in Section 3.

Acknowledgments

AWH was funded by the U.S. Army Research Office under grant W9111NF-05-1-0294, the European Commission under Marie Curie grants ASTQIT (FP6-022194) and QAP (IST-2005-15848), and the U.K. Engineering and Physical Science Research Council through “QIP IRC.”

References

- [Aar03] Scott Aaronson. Quantum lower bound for recursive Fourier sampling. *Quantum Information and Computation*, 3(2):165–174, 2003. arXiv:quant-ph/0209060.
- [AJL06] Dorit Aharonov, Vaughan Jones, and Zeph Landau. A polynomial quantum algorithm for approximating the jones polynomial. In *STOC '06: Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 427–436, New York, NY, USA, 2006. ACM Press. arXiv:quant-ph/0511096.
- [BCvD05] Dave Bacon, Andrew M. Childs, and Wim van Dam. From optimal measurement to efficient quantum algorithms for the hidden subgroup problem over semidirect product groups. In *FOCS '05: 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 469–478, 2005. arXiv:quant-ph/0504083.
- [Bea97] Robert Beals. Quantum computation of Fourier transforms over symmetric groups. In *STOC '97: Proceedings of the Twenty-Ninth Annual ACM Symposium on Theory of Computing*, pages 48–53, El Paso, Texas, 4–6 May 1997. ACM Press.
- [Ber97] Bonnie Berger. The fourth moment method. *Siam J. Comp.*, 26(4):1188–1207, August 1997.
- [BV97] Ethan Bernstein and Umesh Vazirani. Quantum complexity theory. *Siam J. Comp.*, 26(5):1411–1473, Oct 1997.
- [DOP07] Oscar C.O. Dahlsten, Roberto Oliveira, and Martin B. Plenio. Emergence of typical entanglement in two-party random processes. *J. Phys. A*, 40:8081–8108, 2007. arXiv:quant-ph/0701125.
- [DS96] P. Diaconis and L. Saloff-Coste. Logarithmic Sobolev inequalities for finite Markov chains. *Ann. Appl. Probab.*, 6(3):695–750, 1996.
- [FIM⁺03] Katalin Friedl, Gabor Ivanyos, Frederic Magniez, Miklos Santha, and Pranab Sen. Hidden translation and orbit coset in quantum computing. In *STOC '03: Proceedings of the Thirty-Fifth Annual ACM Symposium on Theory of Computing*, pages 1–9, San Diego, CA, 2003. ACM Press.
- [FvdG99] Chris A. Fuchs and Jeroen van de Graaf. Cryptographic distinguishability measures for quantum mechanical states. *IEEE Trans. Inf. Th.*, 45(4):1216–1227, 1999. arXiv:quant-ph/9712042.
- [Hal07] Sean Hallgren. Polynomial-time quantum algorithms for Pell’s equation and the principal ideal problem. *Journal of the ACM*, 54(1):1–19, 2007.
- [HL08] Aram W. Harrow and Richard Low. Random circuits are approximate 2-designs. arXiv:0802.1919, 2008.
- [HMR⁺06] Sean Hallgren, Cristopher Moore, Martin Rötteler, Alexander Russell, and Pranab Sen. Limitations of quantum coset states for graph isomorphism. In *STOC '06: Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 604–617, New York, NY, USA, 2006. ACM Press. arXiv:quant-ph/0511148, arXiv:quant-ph/0612089.
- [IMS01] Gábor Ivanyos, Frédéric Magniez, and Miklos Santha. Efficient quantum algorithms for some instances of the non-abelian hidden subgroup problem. In *Proceedings of the Thirteenth Annual ACM Symposium on Parallel Algorithms and Architectures*, pages 263–270, Heraklion, Crete Island, Greece, 2001.
- [PSW06] Sandu Popescu, Anthony J. Short, and Andreas Winter. The foundations of statistical mechanics from entanglement: Individual states vs. averages. *Nature*, 2:754–758, 2006. arXiv:quant-ph/0511225.
- [Sho97] Peter W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *Siam J. Comp.*, 26(5):1484–1509, October 1997. arXiv:quant-ph/9508027.
- [Sim97] Daniel R. Simon. On the power of quantum computation. *Siam J. Comp.*, 26(5):1474–1483, October 1997.
- [vDHI03] Wim van Dam, Sean Hallgren, and Lawrence Ip. Quantum algorithms for some hidden shift problems. In *SODA '03: Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, Baltimore, MD, 2003.
- [Wat01] John Watrous. Quantum algorithms for solvable groups. In *STOC '01: Proceedings of the Thirty-Third Annual ACM Symposium on Theory of Computing*, pages 60–67, Crete, Greece, 2001. ACM Press. arXiv:quant-ph/0011023.

A Most circuits are dispersing

In this Appendix we prove Theorem 12.

Suppose we start in a computational basis state $|a\rangle$, and after t steps of a random circuit (described in Section 6), we have the state $|\psi_t\rangle$. Let U^\dagger denote the circuit we have applied, and let ψ_t denote $|\psi_t\rangle\langle\psi_t|$, so that $\psi_t = U^\dagger |a\rangle\langle a| U$. For $p \in \{0, 1, 2, 3\}^n$ let σ_p denote the tensor product of Pauli matrices $\sigma_{p_1} \otimes \cdots \otimes \sigma_{p_n}$,

where $\{\sigma_0, \sigma_1, \sigma_2, \sigma_3\}$ are the usual single-qubit Pauli matrices $\{I, \sigma_z, \sigma_x, \sigma_y\}$. Then we can expand ψ_t in the Pauli basis (following [DOP07]) as

$$\psi_t = 2^{-\frac{n}{2}} \sum_p \gamma_t(p) \sigma_p,$$

where $\gamma_t(p) = 2^{-\frac{n}{2}} \text{tr } \psi_t \sigma_p$. The advantage of this approach is that each $\gamma_t(p)$ is real and $\sum_p \gamma_t^2(p) = 1$, so we can think of $\{\gamma_t^2(p)\}$ as a probability distribution on $p \in \{0, 1, 2, 3\}^n$.

Indeed, by an argument similar to that in [DOP07], we can show how $\{\mathbf{E}_C \gamma_t^2(p)\}$ evolves in a way that can be described as a Markov chain on $\{0, 1, 2, 3\}^n$. (Here the expectation is taken over the choice of random quantum circuit.)

Lemma 13. *Random quantum circuits are such that $\{\mathbf{E}_C \gamma_t^2(p)\}$ evolve according to the following Markov chain on $\{0, 1, 2, 3\}^n$:*

- *Select $i \neq j$ randomly from $[n]$.*
- *If $p_i = p_j = 0$ then do nothing.*
- *Otherwise set (p_i, p_j) to a random element of $\{0, 1, 2, 3\}^2 \setminus \{(0, 0)\}$.*

Furthermore, this Markov chain satisfies the following properties.

1. *This Markov chain is irreducible and ergodic once we delete the isolated vertex 0^n .*
2. *Its stationary distribution (when starting with any physical state) has $\gamma^2(0^n) = 2^{-n}$, but otherwise is uniform on $p \neq 0^n$, i.e. $\gamma^2(p) = 4^{-n}/(1 + 2^{-n})$ for all $p \in \{0, 1, 2, 3\}^n$.*
3. *Its spectral gap is $\geq \Omega(1/n^2)$.*
4. *There exists a constant C such that if $t \geq Cn^3$ and if the initial state is a computational basis state then $\mathbf{E}_C \gamma_t^2(p) \leq 4^{-n}$ for all $p \neq 0^n$.*

Before proving the Lemma, we show how it implies Theorem 12. Fix an input $|a\rangle$, apply t random two-qubit unitaries as described above to yield the state $|\psi_t\rangle$, and define $Q_t := \sum_x |\langle x | \psi_t \rangle|^4$. Expanding ψ_t in terms of σ_p , we obtain

$$Q_t = 2^{-n} \sum_x \langle x | \sum_{p \in \{0,1,2,3\}^n} \gamma_t(p) \sigma_p | x \rangle \langle x | \sum_{q \in \{0,1,2,3\}^n} \gamma_t(q) \sigma_q | x \rangle.$$

Now $\langle x | \sigma_p | x \rangle$ will be zero if p contains any 2's or 3's, since each of these lead to bit flips. So we can restrict our sum to p 's and q 's that are strings of 0's and 1's (corresponding to I and σ_z). Moreover, if p is such a string, then $\langle x | \sigma_p | x \rangle = (-1)^{p \cdot x}$. Thus

$$Q_t = 2^{-n} \sum_{p \in \{0,1\}^n, q \in \{0,1\}^n} \gamma_t(p) \gamma_t(q) \sum_x (-1)^{(p+q) \cdot x} = \sum_{p \in \{0,1\}^n} \gamma_t^2(p).$$

To bound this sum, we use the last part of Lemma 13 together with $\gamma_t^2(0^n) = 2^{-n}$ to find

$$\mathbf{E}_C Q_t \leq 2 \cdot 2^{-n}.$$

Thus, Markov's inequality implies that

$$\Pr_C \left(Q_t \geq \frac{2^{-n}}{\beta^2} \right) \leq 2\beta^2.$$

Now consider the event that $Q_t \leq 2^{-n}/\beta^2$. We will use this to lower-bound $\sum_x |\langle x|\psi_t \rangle|$. To do so, we define a random variable Y to equal $|\langle x|\psi_t \rangle|$ for a uniformly random choice of $x \in \{0,1\}^n$. Then $\mathbf{E}_x Y^2 = 2^{-n}$, $\mathbf{E}_x Y^4 \leq \beta^{-2} \cdot 4^{-n}$, and by (11), $\mathbf{E}_x Y = \mathbf{E}_x |Y| \geq 2^{-n/2}\beta$. Thus $\sum_x |\langle x|\psi_t \rangle| \geq \beta 2^{n/2}$.

Putting this together, we see that for any fixed input $|a\rangle$, and for all but a $2\beta^2$ -fraction of (sufficiently long) random circuits,

$$\sum_{x \in \{0,1\}^n} |\langle x|U^\dagger|a \rangle| \geq \beta 2^{\frac{n}{2}}.$$

Say that a pair (U, a) is *bad* when this does not hold. So for any a , the probability over $U \in \mathcal{C}$ that (U, a) is bad is $\leq 2\beta^2$. Thus Markov's inequality implies that

$$\Pr_{U \in \mathcal{C}} \left[\Pr_a [(U, a) \text{ is bad}] \geq 1 - 2^{-(1-\alpha)n} \right] \leq \frac{2\beta^2}{1 - 2^{-(1-\alpha)n}}.$$

Turning this around, we conclude that a random circuit $U \in \mathcal{C}$ is (α, β) -dispersing (meaning that (U, a) is good for $\geq 2^{\alpha n}$ values of a) with probability $\geq 1 - 2\beta^2/(1 - 2^{-(1-\alpha)n})$. Thus, we can obtain a quantum/classical separation from almost any quantum circuit with uniformly bounded parameters for both the quantum upper bound and the classical lower bound.

Recent work[HL08] has improved the analysis of the Markov chain to show that the gap is $\Omega(1/n)$ and hence that circuits of length $O(n^2)$ are generically dispersing.

Proof of Lemma 13: The reduction of the quantum random circuit to a classical Markov chain is due to [DOP07], but we will present an alternate, shorter proof in order to have a self-contained presentation.

First, we show that $\gamma_t(p)^2$ follows a Markov chain. Recall that $\psi_t = 2^{-n/2} \sum_p \gamma_t(p) \sigma_p$ and let W denote the random unitary applied at time $t+1$. Then $\psi_{t+1} = W\psi_t W^\dagger$. Since W acts only on two qubits, we can assume for the purpose of this analysis that $n = 2$, so $W \in \mathcal{U}_4$. Then

$$\gamma_{t+1}(p) = \frac{1}{2} \text{tr } \sigma_p W \psi_t W^\dagger = \frac{1}{4} \sum_{q \in \{0,1,2,3\}^2} \gamma_t(q) \text{tr } \sigma_p W \sigma_q W^\dagger,$$

and

$$\gamma_{t+1}^2(p) = \frac{1}{16} \sum_{q, q' \in \{0,1,2,3\}^2} \gamma_t(q) \gamma_t(q') \text{tr } \sigma_p W \sigma_q W^\dagger \text{tr } \sigma_p W \sigma_{q'} W^\dagger \quad (12)$$

$$= \sum_{q, q' \in \{0,1,2,3\}^2} \gamma_t(q) \gamma_t(q') \langle p, p | \text{ad}_W^{\otimes 2} | q, q' \rangle. \quad (13)$$

Here ad_W is an operator on \mathbb{C}^{16} defined by $\langle p | \text{ad}_W | q \rangle := \text{tr } \sigma_p W \sigma_q W^\dagger / 4$ for $p, q \in \{0,1,2,3\}^2$. When we take the expectation over random choices of W , we obtain

$$\mathbf{E} \gamma_{t+1}^2(p) = \sum_{q, q' \in \{0,1,2,3\}^2} \gamma_t(q) \gamma_t(q') \langle p, p | \left(\mathbf{E}_W \text{ad}_W^{\otimes 2} \right) | q, q' \rangle. \quad (14)$$

Define

$$|\xi\rangle = \frac{1}{\sqrt{15}} \sum_{p \in \{0,1,2,3\}^2 \setminus \{(0,0)\}} |p\rangle |p\rangle. \quad (15)$$

We claim that $\mathbf{E}_W \text{ad}_W^{\otimes 2} = |00\rangle\langle 00| + |\xi\rangle\langle \xi|$. Since $\langle p, p | \xi \rangle \langle \xi | q, q' \rangle$ is equal to $1/15$ when $p \neq 00$ and $q = q' \neq 00$, and zero otherwise, we will have

$$\mathbf{E}_W \langle p, p | \text{ad}_W^{\otimes 2} | q, q' \rangle = \begin{cases} 1 & \text{if } p = q = q' = 00 \\ \frac{1}{15} & \text{if } p \neq 00 \text{ and } q = q' \neq 00 \\ 0 & \text{otherwise} \end{cases},$$

as claimed in the lemma.

To prove that $\mathbf{E}_W \text{ad}_W^{\otimes 2} = |00\rangle\langle 00| + |\xi\rangle\langle \xi|$, we will use representation theory. Schur's Lemma implies that $\mathbf{E}_W \text{ad}_W^{\otimes 2}$ is a projector onto the invariant subspace of $\text{ad}_W^{\otimes 2}$. For any integers $\lambda_1 \geq \dots \geq \lambda_d$, we have an irrep of \mathcal{U}_d , which we call \mathcal{Q}_λ^d . The conjugate irrep of \mathcal{Q}_λ^d , obtained by taking the complex conjugate of each representation matrix, is given by $(\mathcal{Q}_\lambda^d)^* \cong \mathcal{Q}_{\lambda'}^d$, where $\lambda' = (-\lambda_d, -\lambda_{d-1}, \dots, -\lambda_1)$.

The simplest non-trivial \mathcal{U}_d representation is the defining representation $\mathcal{Q}_{(1)}^d \cong \mathbb{C}^d$, where $U \in \mathcal{U}_d$ is mapped to itself. Here we have dropped trailing zeros, so (1) is equivalent to $(1, 0, 0, 0)$. Now observe that $W \rightarrow \text{ad}_W$ is a representation of \mathcal{U}_4 that is equivalent to $\mathcal{Q}_{(1)}^4 \otimes (\mathcal{Q}_{(1)}^4)^* \cong \mathcal{Q}_{(1)}^4 \otimes \mathcal{Q}_{(0,0,0,-1)}^4$. Now we apply the Clebsch-Gordan decomposition of the adjoint representation $\mathcal{Q}_{(1)}^4 \otimes \mathcal{Q}_{(0,0,0,-1)}^4$ to obtain

$$\mathcal{Q}_{(1)}^4 \otimes \mathcal{Q}_{(0,0,0,-1)}^4 \cong \mathcal{Q}_{(0)}^4 \oplus \mathcal{Q}_{(1,0,0,-1)}^4,$$

meaning the direct sum of a trivial irrep $\mathcal{Q}_{(0)}^4$ and a 15-dimensional irrep $\mathcal{Q}_{(1,0,0,-1)}^4$. To find the invariant subspace of $\text{ad}_W^{\otimes 2}$, we will need to find the invariant subspace of

$$(\mathcal{Q}_{(0)}^4 \oplus \mathcal{Q}_{(1,0,0,-1)}^4)^{\otimes 2} = (\mathcal{Q}_{(0)}^4)^{\otimes 2} \oplus (\mathcal{Q}_{(0)}^4 \otimes \mathcal{Q}_{(1,0,0,-1)}^4) \oplus (\mathcal{Q}_{(1,0,0,-1)}^4 \otimes \mathcal{Q}_{(0)}^4) \oplus (\mathcal{Q}_{(1,0,0,-1)}^4)^{\otimes 2}.$$

Since $\mathcal{Q}_{(0)}^4$ is the trivial representation, $\mathcal{Q}_{(0)}^4 \otimes \mathcal{Q}_{(0)}^4$ is trivial as well, and the projector onto it is given by $|00\rangle\langle 00|$. Next $\mathcal{Q}_{(0)}^4 \otimes \mathcal{Q}_{(1,0,0,-1)}^4 \cong \mathcal{Q}_{(1,0,0,-1)}^4$, and thus has no invariant subspace. Similarly for $\mathcal{Q}_{(1,0,0,-1)}^4 \otimes \mathcal{Q}_{(0)}^4$ has no invariant subspace. $\mathcal{Q}_{(1,0,0,-1)}^4$ is self-dual, and thus by Schur's Lemma, $\mathcal{Q}_{(1,0,0,-1)}^4 \otimes \mathcal{Q}_{(1,0,0,-1)}^4$ has a one-dimensional invariant subspace. To determine this subspace we observe that in the basis $\{|p\rangle\}_{p \in \{0,1,2,3\}^2 \setminus \{(0,0)\}}$, the representation matrices of $\mathcal{Q}_{(1,0,0,-1)}^4$ are real. Also $(A \otimes I)|\xi\rangle = (I \otimes A^T)|\xi\rangle$ for any matrix A and for $|\xi\rangle$ defined in (15). Together this means that $|\xi\rangle$ is an invariant vector in $\mathcal{Q}_{(1,0,0,-1)}^4 \otimes \mathcal{Q}_{(1,0,0,-1)}^4$, and in fact spans its invariant subspace. We conclude that $\mathbf{E}_W \text{ad}_W^{\otimes 2} = |00\rangle\langle 00| + |\xi\rangle\langle \xi|$.

We now turn to the analysis of the classical Markov chain. Similar arguments were used in [DOP07] and a tighter analysis is forthcoming in [HL08]. First, we claim that the Markov chain is ergodic and irreducible outside the vertex 0^n . Irreducibility follows from the fact that every nonzero vertex is connected to 1^n , while the presence of self-loops implies ergodicity. We can verify the stationary distribution from the detailed balance condition using a short calculation.

For the gap we will use the comparison method of Diaconis and Saloff-Coste [DS96]. Consider a Markov chain that picks two random sites and replaces them each with random numbers from $\{0, 1, 2, 3\}$ subject only to avoiding the state 0^n . It follows from [DS96, Thm 3.2] that this chain has gap $\geq 2/n$, and applying the comparison method ([DS96, Thm 3.3]) to the Markov chain described in Lemma 13 yields a lower bound of $1/n^2$ for its gap.

To prove the final claim of Lemma 13, we want to choose C sufficiently large so that

$$|\mathbf{E} \gamma_t^2(p) - 4^{-n}/(1 + 2^{-n})| \leq 2^{-4n} \quad (16)$$

whenever $t \geq Cn^3$. Note that any computational basis state has overlap $\exp(-O(n))$ with the stationary distribution. Since the gap is $\Omega(1/n^2)$, we can then use standard bounds on the mixing time of Markov chains (e.g. [DS96, Lemma 2.8]) to show that (16) holds when $t \geq Cn^3$ for some constant C .

B Proof of Lemma 10

We want to bound

$$\max_{\theta} \left| \sum_{k=1}^d \theta_k x_k \right| = \max_{\phi} \max_{\theta} \text{Re} \left(e^{i\phi} \sum_{k=1}^d \theta_k x_k \right).$$

Here ϕ is maximized over $[0, 2\pi]$, and $\text{Re}(z)$ refers to the real part of a complex number z . We can now move the maximization over θ inside the sum to obtain

$$\max_{\phi} \max_{\theta} \text{Re} \left(e^{i\phi} \sum_{k=1}^d \theta_k x_k \right) = \max_{\phi} \sum_{k=1}^d \max_{\theta_k \in \{\pm 1\}} \text{Re} \left(e^{i\phi} \theta_k x_k \right) \quad (17)$$

$$= \max_{\phi} \sum_{k=1}^d \left| \text{Re}(e^{i\phi} x_k) \right| \geq \mathbf{E}_{\phi} \sum_{k=1}^d \left| \text{Re}(e^{i\phi} x_k) \right| \quad (18)$$

$$= \mathbf{E}_{\phi} \sum_{k=1}^d |x_k| \cdot |\cos \phi| = \frac{2}{\pi} \sum_{k=1}^d |x_k|. \quad (19)$$

\mathbf{E}_{ϕ} indicates the expectation over ϕ chosen uniformly at random from $[0, 2\pi]$, and the second to last equality uses the rotational invariance of the distribution of ϕ . \square